

Listing of Claims

This Listing of Claims will replace all prior versions, and listings, of claims in the application.

1-20. (Cancelled)

21. (Previously Presented) A method of managing a cluster of networked resources and resource groups using rule-based constraints in a scalable clustering environment, the method comprising the steps of:

building a globally optimal cluster configuration of said networked resources in accordance with said rule-based constraints and a current state of said resources, including identifying for each of the resources and resource groups an availability and quality of service, which are determined by dependencies among the resources and resource groups, resource equivalency, user preferences, constraints on the resources and network policies[[,]];

bringing said cluster of networked resources on-line in a systematic manner, given the current state of each of the resources and resource groups, and their dependencies, user preferences, constraints on the resources, and network policies,

with said cluster of networked resources on-line, determining dynamic dependencies of and configuration information about said cluster of networked resources (i) statically at said step of building and said step of bringing said cluster of networked services online and (ii) dynamically during cluster operation in accordance with said rule-based constraints,

supporting startup, operation and shutdown of said cluster of networked resources according to current policies, and system events, and said rule-based constraints;

separating said dependencies among resources and resource groups, user preferences, constraints among the resources, system events, and current policies into (i) a first, static rules based group and (ii) a second, dynamically changing events based group, wherein said first group captures the static resources, including, for each resource, a type and quality of the supporting resources needed to enable said each resource, wherein said step of separating is implemented according to said rule-based constraints; and

combining said first and second groups in a systematic manner only when needed to build the said globally optimal cluster configuration, and only when needed during operation to modify and realign the current state of said cluster to said globally optimal cluster configuration, or an alternative globally optimal in view of said current policies, said system events and said rule-based constraints.

22. (Previously Presented) A method according to Claim 21, wherein said step of combining further comprising:

continuously monitoring system events and comparing the current cluster state with the globally optimal cluster configuration, and upon detecting a discrepancy between said current cluster state and said globally optimal cluster configuration, realigning said cluster of networked resources, including issuing commands to the networked resources comprising the current cluster state to bring about the realigning;

providing a group of cluster resources, including:

- i) a persistent cluster registry to store and retrieve the configuration of the cluster of the networked resources,
- ii) topology resources for detecting node and communication adapter failures within said cluster of networked resources,
- iii) messaging for selected communications between a central resource and all other resources comprising the cluster of networked resources, and
- iv) a group resources facility for electing one of the resources as the central resource at cluster initialization, and upon determining that an existing central resource is unable to provide the resources thereof,

delivering notification of system events to a coordinator to process said system events in accordance with said rule-based constraints to arrive at a response to said system events;

translating the response into commands to the resources that realign the availability, quality of service and related dependencies to execute said commands by a resource manager associated with each of said resources, including issuance of the commands in a partial order sequence where necessary, in view of said each resource's dependencies; and

not sending out a "next" command until the central resource detects a positive outcome of the current command or commands that the execution of said next command depends on.

23. (Previously Presented) A method according to Claim 22, wherein:

said coordinator, by said rule-based constraints, ensures that the current cluster state is realigned with said globally optimal cluster configuration, or an alternative globally optimal cluster configuration in response to system events in said cluster, wherein all events and command

feedback from said cluster in response to a current or next command is directed to said coordinator.

24. (Previously Presented) A method according to Claim 21, further comprising:

providing an optimizer module for computing said globally optimal cluster configuration, or an alternative globally optimal cluster configuration based on said rule-based constraints and to realign the current state of said cluster when needed in response to system events;

using the optimizer module for realigning the current configuration state in accordance with said rule-based constraints whenever an objective value calculated in view of the current configuration state is below a certain value calculated for the globally optimal cluster configuration, including feeding back to the optimizer module an artificially generated event that forces the optimizer to realign the current network configuration to a cluster configuration approaching or equivalent to said global optimal cluster configuration and said alternative global cluster configuration;

providing the optimizer module with a snapshot of a current state of the cluster of networked resources;

wherein the optimizer, given said snapshot, calculates and proposes an approximately globally optimal cluster configuration that takes into account said current state of the cluster and long-term objectives defined for the cluster in accordance with said rule-based constraints.

25. (Previously Presented) A system for managing a cluster of networked resources and resource groups using rule-based constraints in a scalable clustering environment, comprising apparatus for:

building a globally optimal cluster configuration of the networked resources, in accordance with said rule-based constraints and a current state of said resources and including identifying each of the resources with an availability and quality of service, which are determined by dependencies among the resources and resource groups, resource equivalency, user preferences, constraints on the resources, and network policies,

bringing said cluster of networked resources on-line in a systematic manner, given a current state of each of the resources and resource groups, and said dependencies, user preferences, constraints on the resources and policies, in accordance with said rule-based constraints,

with said cluster of networked resources on-line, determining dynamic dependencies of configuration information about said cluster of networked resources (i) statically at said step of building and said step of bringing said cluster of networked resources online and (ii) dynamically during cluster operation,

supporting a startup, operation and shutdown of said cluster of networked resources according to current policies, and system events, in accordance with said rule-based constraints,

separating the said dependencies among resources, resource groups, user preferences, constraints among the resources, system events, equivalencies among said resources and said resource groups and current policies into (i) a first, static rules based group and (ii) a second, dynamically changing events based group and (ii) a second, dynamically changing events based

group, wherein said first group captures the static resources, including, for each resource and resource group, the type and quality of the supporting resources, resource groups and their equivalencies needed to enable said each resource and resource groups, and

combining said first, static and second, dynamically changing events based groups in a systematic manner only when needed to build the said globally optimal cluster configuration at startup, and only when needed during operation to modify and realign the current state of said cluster to said globally optimal cluster configuration in view of said current policies, said system events, said equivalencies and said rule-based constraints.

26. (Previously Presented) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps of managing a cluster of networked resources, including resource groups, using rule-based constraints in a scalable clustering environment, said method steps comprising:

building a globally optimal cluster configuration of said networked resources, in accordance with said rule-based constraints and a current state of said resources, including identifying each of the resources with an availability and quality of service, which are determined by dependencies among the resources and resource groups, resource equivalency, user preferences, constraints on the resources and network policies,

bringing said cluster of networked resources on-line in a systematic manner, given a current state of each of the resources and resource groups, and said dependencies, user preferences, constraints on the resources, and network policies,

with said cluster of networked resources on-line, determining dynamic dependencies of configuration information about said cluster of networked resources and resource groups (i) statically at said step of bringing said cluster of networked resources on-line and (ii) dynamically during cluster operation, in accordance with said rule-based constraints,

supporting startup, operation and shutdown of said cluster of networked resources according to current policies, and system events, and said rule-based constraints,

separating said dependencies among resources and resource groups, user preferences, constraints among the resources, system events, equivalencies and current policies into (i) a first, static rules based group and (ii) a second, dynamically changing events based group, wherein said first group captures the static resources, including, for each resource and resource groups, the type and quality of the supporting resources, and equivalencies of said resources and said resource groups needed to enable said each resource, and said resource groups, and

combining said first and second groups only when needed to build the said globally optimal cluster configuration, and only when needed during operation to modify and realign the current state of said cluster to said globally optimal cluster configuration in view of said current policies, said system events, said equivalencies and said rules-based constraints, to achieve a desired level of automation in the coordination and mapping of networked resources.

27. (Previously Presented) A method according to Claim 24, wherein:

the providing step includes providing a preprocessor module and a postprocessor module;

wherein the preprocessor module includes a preprocessor entry queue, the optimizer module includes an optimizer input queue, and the postprocessor module includes a postprocessor input queue;

creating a preprocessor task to implement a decision to reallocate a resource or a resource group, wherein said preprocessor task is provided to the entry queue of the preprocessor module; and

wherein said preprocessor task is an object having an entry method that, when invoked, results in the preprocessor task being executed, and execution of the preprocessor task results in either a postprocessor task being provided in the postprocessor input queue, an optimizer task being provided in the optimizer input queue, or both;

scheduling the postprocessor task by an invocation of the entry method associated with the postprocessor task;

scheduling the optimizer task by an invocation of the entry method associated with the optimizer task; and

executing the optimizer task results in a postprocessor task provided in the postprocessor input queue.

28. (Withdrawn) A system for managing a dynamic scalable cluster of heterogeneous networked resources, and networked resource groups, by arranging the resources and resource groups in an optimal cluster configuration for servicing dynamic scalable resource needs of a network of computer systems in accordance with a set of rule-based constraints, wherein each said resource and resource groups includes a resource manager and a set of attributes, and is

allocated and reallocated by the system in view of its dependencies to other resources, equivalencies to other resources and a specified load-based policy for responding to event-driven changes in resource needs, and/or changes in said attributes defining each of said heterogeneous resources, the system comprising:

- a cluster registry to store and recall attributes and resource manager data that define each said heterogeneous resource and resource group;

- a mechanism for detecting differences between current resource attributes defining each of said heterogeneous resources and resource groups, and stored attributes associated with said globally optimal cluster configuration;

- an optimization problem solver that responds to differences between detected current resource attributes and attributes associated with said optimal cluster constraint-based configuration by generating a set of instructions for controlling and realigning said current, detected resource attributes; and

- a coordinator for implementing said set of instructions to control said detected current, resource attributes to operate said cluster according to said optimized cluster constraint-based configuration.

29. (Withdrawn) The system as set forth in claim 28, wherein said attributes include name, type, capacity, priority and state, and wherein said state indicates readiness of each said resources' availability as one of offline, online and failed.

30. (Withdrawn) The system as set forth in claim 28, wherein said attributes include resource availability, equivalencies and quality of service.

31. (Withdrawn) The system as set forth in claim 28, wherein said optimization problem solver operates locally when operating upon a proper subset of resources, and operates globally when operating upon said cluster.